

Self-Healing Computation^{*}

George Saad and Jared Saia

Department of Computer Science, University of New Mexico,
 {saad, saia}@cs.unm.edu

Abstract. In the problem of reliable multiparty computation (RC), there are n parties, each with an individual input, and the parties want to jointly compute a function f over n inputs. The problem is complicated by the fact that an omniscient adversary controls a hidden fraction of the parties.

We describe a self-healing algorithm for this problem. In particular, for a fixed function f , with n parties and m gates, we describe how to perform RC repeatedly as the inputs to f change. Our algorithm maintains the following properties, even when an adversary controls up to $t \leq (\frac{1}{4} - \epsilon)n$ parties, for any constant $\epsilon > 0$. First, our algorithm performs each reliable computation with the following amortized resource costs: $O(m + n \log n)$ messages, $O(m + n \log n)$ computational operations, and $O(\ell)$ latency, where ℓ is the depth of the circuit that computes f . Second, the expected total number of corruptions is $O(t(\log^* m)^2)$, after which the adversarially controlled parties are effectively quarantined so that they cause no more corruptions.

Keywords: Self-Healing Algorithms, Threshold Cryptography, Leader Election

1 Introduction

How can we protect a network against adversarial attack? A traditional approach provides robustness through redundant components. If one component is attacked, the remaining components maintain functionality. Unfortunately, this approach incurs significant resource cost, even when the network is not under attack.

An alternative approach is self-healing, where a network automatically recovers from attacks. Self-healing algorithms expend additional resources only when it is necessary to repair from attacks.

In this paper, we describe self-healing algorithms for the problem of *reliable multiparty computation* (RC). In the RC problem, there are n parties, each with an individual input, and the parties want to jointly compute a function f over n inputs. A hidden $1/4$ -fraction of the parties are controlled by an omniscient Byzantine adversary. A party that is controlled by the adversary is said to be *bad*, and the remaining parties are said to be *good*. Our goal is to ensure that all good parties learn the output of f .¹

RC abstracts many problems that may occur in high-performance computing, sensor networks, and peer-to-peer networks. For example, we can use RC to enable performance profiling and system monitoring, compute order statistics, and enable public voting.

Our main result is an algorithm for RC that 1) is asymptotically optimal in terms of total messages and total computational operations; and 2) limits the expected total number of corruptions. Ideally, each bad party would cause $O(1)$ corruptions; in our algorithm, each bad party causes an expected $O((\log^* m)^2)$ corruptions.

^{*} This research is partially supported by NSF Award SATC 1318880.

¹ Note that RC differs from secure multiparty computation (MPC) only in that there is no requirement to keep inputs private.

1.1 Our Model

We assume a static Byzantine adversary that takes over $t \leq (\frac{1}{4} - \epsilon)n$ parties before the algorithm begins, for any constant $\epsilon > 0$. As mentioned previously, parties that are compromised by the adversary are called *bad*, and the remaining parties are *good*. The bad parties may arbitrarily deviate from the protocol, by sending no messages, excessive numbers of messages, incorrect messages, or any combination of these. The good parties follow the protocol. We assume that the adversary knows our protocol, but is unaware of the random bits of the good nodes. We make use of a public key cryptography scheme, and thus assume that the adversary is computationally bounded.

We assume a partially synchronous communication model. Any message sent from one good node to another good node requires at most h time steps to be sent and received, and the value h is known to all nodes. However, we allow the adversary to be *rushing*: the bad nodes receive all messages from good nodes in a round before sending out their own messages. We further assume that each party has a unique ID. We say that party p has a link to party q if p knows q 's ID and can thus directly communicate with node q .

In the reliable multiparty computation problem, we assume that the function f can be implemented with an arithmetic circuit over m gates, where each gate has two inputs and at most two outputs.² For simplicity of presentation, we focus on computing a single function multiple times (with changing inputs). However, we can also compute multiple functions with our algorithm.

1.2 Our Result

We describe an algorithm, *COMPUTE*, to efficiently solve reliable multiparty computation. Our main result is summarized in the following theorem.

Theorem 1. *Assume we have n parties providing inputs to a function f that can be computed by an arithmetic circuit with depth ℓ and containing m gates. Then *COMPUTE* solves RC and has the following properties.*

- In an amortized sense,³ any execution of *COMPUTE* requires the following expected costs:
 - $O(m + n \log n)$ messages sent by all parties,
 - $O(m + n \log n)$ computational operations performed by all parties, and
 - $O(\ell)$ latency.
- The expected total number of times *COMPUTE* returns a corrupted output is $O(t(\log^* m)^2)$.

1.3 Technical Overview

Our algorithms make critical use of quorums and a quorum graph.

Quorums and the Quorum Graph: We define a quorum to be a set of $\Theta(\log n)$ parties, of which at most $1/4$ -fraction are bad. Many results show how to create and maintain a network of quorums [6,11,18,23,7,1,14]. All of these results maintain what we will call a *quorum graph* in which each vertex represents a quorum. The properties of the quorum graph are: 1) each party is in $\Theta(\log n)$ quorums; 2) for any quorum Q , any party in Q can communicate directly to any other party in Q ; and 3) for any quorums Q and Q' that are connected in the quorum graph, any party in Q can communicate directly with any party in Q' and vice versa. Moreover, we assume that for any two parties x and y in a quorum, x knows all quorums that y is in.

² We note that any gate of any fixed in-degree and out-degree can be converted into a fixed number of gates with in-degree 2 and out-degree at most 2.

³ In particular, if we call *COMPUTE* \mathcal{L} times, then the expected total number of messages sent will be $O(\mathcal{L}(m + n \log n) + t(m \log^2 n))$. Since t is fixed, for large \mathcal{L} , the expected number of messages per *COMPUTE* is $O(m + n \log n)$. The result for computational operations is similar.




Fig. 1. Quorum-Based Circuit

Computing with Quorums: We maintain a quorum graph with $m + n$ nodes: m nodes for the gates of the circuit and n nodes for the inputs of the parties. The input nodes are connected to the gates using these inputs, and the gate nodes are connected as in the circuit. Quorums are mapped to nodes in this quorum graph as described above. See Figure 1. Intuitively, the computation is performed from the left to the right, where the input quorums are the leftmost quorums and the output quorum is the rightmost one.

A correct but inefficient way to solve RC is as follows. Each party s_i sends its input to all parties of the appropriate input quorum. Then the computation is performed from left to right. All parties in each quorum compute the appropriate gate operation on their inputs, and send their outputs to all parties in the right neighboring quorums via all-to-all communication. At the next level, all parties in each quorum take the majority of the received messages in order to determine the correct input for their gate. At the end, the parties in the rightmost quorum will compute the correct output of the circuit. They then forward this output back from right to left through the quorum graph using the same all-to-all communication and majority filtering.

Unfortunately, this naive algorithm requires $O(m \log^2 n)$ messages and $O(m \log n)$ computational operations. Our main goal is to remove the logarithmic factors.⁴

Leaders and the CHECK Algorithm: A more efficient approach is for each quorum to have a leader, and for this leader to receive inputs, perform gate computations, and send off the output. Unfortunately, a single bad leader can corrupt the entire computation.

To address this issue, we create *CHECK* (Section 2.3). This algorithm determines if there has been a corruption, and if so, returns at least one pair of parties that are in *conflict*. Informally, we say that a pair of parties are in conflict if they each accuse the other of malicious behavior. In such a situation, we know that at least one party in the pair is bad. Our approach is to mark both parties in each conflicting pair, and then to forbid any marked party from being a leader of a quorum.⁵

The basic idea of *CHECK* is to have multiple rounds where in each round, a new party is selected independently at random from each quorum. We call these parties the *checkers*. For convenience of presentation, we will refer to the leaders as the checkers for round 0. For each round $i \geq 1$, all i checkers at gate g : 1) receive inputs to g from all i checkers at input gates for g ; 2) compute the gate

⁴ We note that such asymptotic improvements can be significant for large networks. For example, if $n = 64,000$, then we would expect our algorithm to reduce message costs by a factor of $\log^2 n = 255$.

⁵ A technical point is that we may need to unmark all parties in a quorum if too many parties in that quorum become marked. However, a potential function argument (Lemma 5) shows that it is still the case that after $O(t)$ markings, all bad parties will be marked.

output for g based on these inputs; and 3) send this output to all i checkers at each output gate for g . If a checker ever receives inconsistent inputs, it calls *UPDATE* (Section 2.4), which will return at least one pair of parties that are in conflict. Unfortunately, waiting until a round where each gate has had at least one good checker would require $O(\log n)$ rounds.

To do better, we use the following approach. Let G be the quorum graph as defined above and let the checkers be selected as above. Call a subgraph of G bad in a given round if all checkers in the nodes of that subgraph are bad in that round. Any corruption in the first round must occur in a bad subgraph. Moreover, for i rounds of *CHECK* to fail to find a corruption, there must be nesting levels of bad subgraphs in G in each of those i rounds.

When *CHECK* elects a good checker at a quorum, it is as if it is removing the node associated with the quorum from the quorum graph. Thus, we can view *CHECK* as repeatedly removing nodes from increasingly smaller subgraphs of G until no nodes remain, at which the corruption is detected. A key lemma (Lemma 1) shows that for any directed acyclic graph (DAG) with m nodes and maximum degree 4, when each node is deleted with probability $3/4$, the probability that a connected subgraph of size $\Omega(\log m)$ survives is at most $1/2$. Using this lemma, we can show that *CHECK* requires only $O(\log^* m)$ rounds to detect a corruption with constant probability.⁶

Since *CHECK* requires $O((m + n \log n)(\log^* m)^2)$ messages, we can call it with probability $1/(\log^* m)^2$ and obtain asymptotically optimal resource costs for the RC problem, while incurring an expected $O(t(\log^* m)^2)$ corruptions.

1.4 Related Work

Our results are inspired by recent work on self-healing algorithms. Early work of [8,12,17,24,25] discusses different restoration mechanisms to preserve network performance by adding capacity and rerouting traffic streams in the presence of node or link failures. This work presents mathematical models to determine global optimal restoration paths, and provides methods for capacity optimization of path-restorable networks.

More recent work [4,22,10,9,19,5] considers models where the following process repeats indefinitely: an adversary deletes some nodes in the network, and the algorithm adds edges. The algorithm is constrained to never increase the degree of any node by more than a logarithmic factor from its original degree. In this model, researchers have presented algorithms that ensure the following properties: the network stays connected and the diameter does not increase by much [4,22,10]; the shortest path between any pair of nodes does not increase by much [9]; expansion properties of the network are approximately preserved [19]; and keeping network backbones densely connected [5].

This paper particularly builds on [16]. That paper describes self-healing algorithms that provide reliable communication, with a minimum of corruptions, even when a Byzantine adversary can take over a constant fraction of the nodes in a network. While our attack model is similar to [16], reliable *computation* is more challenging than reliable communication, and hence this paper requires a significantly different technical approach. Additionally, we improve the fraction of bad parties that can be tolerated from $1/8$ to $1/4$.

Reliable multiparty computation (RC) is closely related to the problem of secure multiparty computation (MPC) which has been studied extensively for several decades (see e.g. [26,2,3,21] or the recent book [20]). RC is simpler than MPC in that it does not require inputs of the parties to remain private. Our algorithm for RC is significantly more efficient than current algorithms for MPC, which require at least polylogarithmic blowup in communication and computational costs in order to tolerate a Byzantine adversary. We reduce these costs through our self-healing approach, which expends additional resources only when corruptions occur, and is able to “quarantine” bad parties after $O(t(\log^* m)^2)$ corruptions.

⁶ This probability can be made arbitrarily close to 1 by adjusting the hidden constant in the $O(\log^* m)$ rounds.

1.5 Organization of Paper

The rest of this paper is organized as follows. In Section 2, we describe our algorithms. The analysis of our algorithms is shown in Section 3. Finally, we conclude and describe problems for future work in Section 4.

2 Our Algorithms

In this section, we describe our algorithms: *COMPUTE*, *COMPUTE-CIRCUIT*, *CHECK* and *UPDATE*.

Our algorithms aim at detecting corruptions and marking the bad parties. Parties that are marked are not allowed to participate in the computation process. Initially, all parties are unmarked.

Recall that there are n parties, each provides an input to an input quorum, Q_i , for $1 \leq i \leq n$; and then the computation is performed through m quorums Q_j , for $n + 1 \leq j \leq m + n$. The result is produced at an output quorum Q_{m+n} .

Before discussing our main *COMPUTE* algorithm, we describe that when a party x broadcasts a message msg , signed by a quorum Q , to a set of parties S , it calls *BROADCAST*(msg, Q, S).

2.1 BROADCAST

In *BROADCAST* (Algorithm 1), we use threshold cryptography to avoid the overhead of Byzantine Agreement. In a (η, η') -threshold cryptographic scheme, a private key is distributed among η parties in such a way that 1) any subset of more than η' parties can jointly reassemble the key; and 2) no subset of at most η' parties can recover the key. The private key can be distributed using a *Distributed Key Generation* (DKG) protocol[13]. DKG generates the public/private key shares of all parties in every quorum. The public key of each quorum is known to all parties in the quorum, and to all parties in all neighboring quorums in the circuit.

Algorithm 1 *BROADCAST*(msg, Q, S) ▷ A party x calls this procedure in order to send a message msg , signed by quorum Q , to a set of parties S .

- 1: Party x sends message msg to all parties in Q .
 - 2: Each party in Q signs msg by its private key share to obtain its message share.
 - 3: Each party in Q sends its message share back to party x .
 - 4: Party x interpolates at least $\frac{3|Q|}{4}$ message shares to obtain a signed-message of Q .
 - 5: Party x sends this signed-message to all parties in S .
-

In particular, we use a $(|Q|, \frac{3|Q|}{4} - 1)$ -threshold scheme, where $|Q|$ is the quorum size. A party x calls *BROADCAST* in order to send a message msg to all parties in S so that: 1) at least 3/4-fraction of the parties in quorum Q have received the same message msg ; 2) they agree upon the content of msg ; and 3) they give a permission to x to broadcast this message.

Any call to *BROADCAST* has $O(\log n + |S|)$ messages, $O(\log n)$ computational operations for signing the message msg by $O(\log n)$ parties in Q and $O(1)$ latency.

2.2 COMPUTE

Now we describe our main algorithm, *COMPUTE* (Algorithm 2), which calls *COMPUTE-CIRCUIT* (Algorithm 3). In *COMPUTE-CIRCUIT*, the n parties broadcast their inputs to the input quorums. The input quorums forward these inputs to a circuit of m leaders in order to perform the computation and provide the result to the output quorum. Then this result is sent through the same circuit of leaders back to all parties. Recall that a leader of a quorum, is a party in this quorum, that is: 1)

a representative of all parties in the quorum; and 2) it is known to all parties in the quorum and neighboring quorums. We assume that all parties provide their inputs to the circuit in the same round.

Algorithm 2 COMPUTE ▷ performs a reliable computation and sends the result reliably to all parties.

- 1: COMPUTE-CIRCUIT ▷ computes and sends back the result through a circuit of leaders.
 - 2: TRIGGER-CHECK ▷ The output quorum triggers *CHECK* with probability $1/(\log^* m)^2$.
-

In the presence of an adversary, *COMPUTE-CIRCUIT* is vulnerable to corruptions. Thus, *COMPUTE* calls *TRIGGER-CHECK* (Algorithm 4), in which the parties of the output quorum decide together, to trigger *CHECK* (Algorithm 5) with probability $1/(\log^* m)^2$, using secure multiparty computation (MPC) [21]. *CHECK* is triggered in order to detect if a computation was corrupted in the last call to *COMPUTE-CIRCUIT*, with probability at least $1/2$.

Algorithm 3 COMPUTE-CIRCUIT ▷ performs a computation through a circuit of leaders producing a result at the output quorum; and it sends back result through same circuit to all parties.

- 1: **for** $i = 1, \dots, n$ **do** ▷ provides the inputs to the circuit
 - 2: Party s_i broadcasts its input a_i to all parties in quorum Q_i .
 - 3: All parties in Q_i send a_i to the leaders of the right neighboring quorums of Q_i .
 - 4: **end for**
 - 5: **for** $i = n + 1, \dots, m + n - 1$ **do** ▷ performs the computation.
 - 6: **for all** $j : i < j \leq m + n$ and $(Q_i, Q_j) \in \text{Circuit}$ **do**
 - 7: **if** leader $q_i \in Q_i$ receives all its inputs **then**
 - 8: q_i performs an operation on its inputs producing an output, b_i .
 - 9: q_i sends b_i to leader $q_j \in Q_j$.
 - 10: **end if**
 - 11: **end for**
 - 12: **end for**
 - 13: **if** leader $q_{m+n} \in Q_{m+n}$ receives all its inputs **then** ▷ produces and broadcasts result to output quorum.
 - 14: q_{m+n} performs an operation on its inputs producing an output, b_{m+n} .
 - 15: q_{m+n} broadcasts b_{m+n} to all parties in Q_{m+n} .
 - 16: **end if**
 - 17: **for** $i = m + n, \dots, n + 2$ **do** ▷ sends back the result to the leftmost leaders.
 - 18: **for all** $j : n + 1 \leq j < i$ and $(Q_j, Q_i) \in \text{Circuit}$ **do**
 - 19: Leader $q_i \in Q_i$ sends b_{m+n} to leader $q_j \in Q_j$.
 - 20: **end for**
 - 21: **end for**
 - 22: **for** $i = 1, \dots, n$ **do** ▷ sends result to all parties after broadcasting it to the input quorums.
 - 23: The leaders of the right neighboring quorums of Q_i broadcast b_{m+n} to all parties in Q_i .
 - 24: All parties in Q_i send b_{m+n} to sender s_i .
 - 25: **end for**
-

Unfortunately, while *CHECK* can determine if a corruption occurred, it does not specify the location where the corruption occurred. Thus, if *CHECK* detects a corruption, *UPDATE* (Algorithm 10) is called. When *UPDATE* is called, it identifies two neighboring quorums Q and Q' in the circuit,

Algorithm 4 TRIGGER-CHECK ▷ The parties of the output quorum Q_{m+n} trigger *CHECK* with probability $1/(\log^* m)^2$.

- 1: Each party in Q_{m+n} chooses an input: a real number uniformly distributed between 0 and 1.
 - 2: The parties of Q_{m+n} perform MPC to find the output, $prob$, which is the sum of all their inputs divided by $|Q_{m+n}|$.
 - 3: **if** $prob \leq 1/(\log^* m)^2$ **then**
 - 4: CHECK
 - 5: **end if**
-

such that at least one pair of parties in these quorums is in conflict and at least one party in this pair is bad. Then quorums Q and Q' mark these parties and notify all other quorums that these parties are in. All quorums in which these parties are notify their neighboring quorums. For each pair of leaders that is in conflict, their quorums elect a new pair of unmarked leaders uniformly at random. If $(1/2 - \gamma)$ -fraction of parties in any quorum have been marked, for any constant $\gamma > 0$, e.g., $\gamma = 0.01$, they are set unmarked in all their quorums and neighboring quorums.

Moreover, we use *BROADCAST* in *COMPUTE-CIRCUIT* and *CHECK* in order to handle any accusation issued in *UPDATE* against the parties that provide the inputs to the input quorums, or those that receive the result in the output quorum.

Our model does not directly consider concurrency. In a real system, concurrent executions of *COMPUTE* that overlap at a single quorum may allow the adversary to achieve multiple corruptions at the cost of a single marked bad party. However, this does not effect correctness, and, in practice, this issue can be avoided by serializing concurrent executions of *COMPUTE*. For simplicity of presentation, we leave the concurrency aspect out of this paper.

2.3 CHECK

In this section, we describe *CHECK* algorithm, which is stated formally as Algorithm 5. In this algorithm, we make use of subquorums, where a subquorum is a subset of unmarked parties in a quorum. Let U_k be the set of all unmarked parties in quorum Q_k , for $1 \leq k \leq m+n$.

Algorithm 5 CHECK ▷ Party r calls *CHECK* to check for corruptions.

- 1: Let subquorums, S_j^1, S_j^2 and S_j^3 , be initially empty, for all $n+1 \leq j \leq m+n$.
- 2: **for** $i \leftarrow 1, \dots, 16 \log^* m$ **do**
- 3: ELECT(Q_{m+n}) ▷ elects a party $r \in Q_{m+n}$.
- 4: Party r constructs R_1, R_2 and R_3 to be three m by m' arrays of random numbers, where m' is the maximum size of any quorum. Note that $R_1[k, k'], R_2[k, k']$ and $R_3[k, k']$ are three uniformly random numbers between 1 and k' , for $1 \leq k \leq m$ and $1 \leq k' \leq m'$.
- 5: REQUEST(R_1, R_2, i) ▷ r requests all senders to recompute.
- 6: RECOMPUTE(R_2, i, r) ▷ recomputes producing the result, b_{m+n} , at r .
- 7: RESEND-RESULT(R_3, b_{m+n}, i) ▷ r sends back b_{m+n} to all parties.
- 8: **end for**

Note that: for each algorithm that is called in *CHECK*, if any party has previously received a message, it verifies this message with all subsequent messages; also if a party receives inconsistent messages or fails to receive an expected message, then it initiates a call to *UPDATE*.

CHECK runs for $O(\log^* m)$ rounds. For each round i , the parties of the output quorum Q_{m+n} elect an unmarked party r from Q_{m+n} to be in charge of the recomputation in round i . This election

Algorithm 6 $ELECT(Q)$ \triangleright Parties in Q elect an unmarked party in Q using MPC.

- 1: Let each party in the set of unmarked parties, $U \subset Q$, is assigned a unique integer from 1 to $|U|$.
 - 2: Each party in Q chooses an input: an integer uniformly distributed between 1 and $|U|$.
 - 3: The parties of Q perform MPC to find the output: the sum of all their inputs modulo $|U|$.
 - 4: The party in U associated with this output number is the elected party.
-

process is stated formally in *ELECT* (Algorithm 6). The elected party r calls *REQUEST* (Algorithm 7) to send a request through a DAG of subquorums, S_j^1 's, to the n senders to recompute. Then the recomputation starts by *RECOMPUTE* (Algorithm 8), in which each sender that receives such request provides its input to redo the computation through a DAG of subquorums, S_j^2 's. Finally, when r receives the result of the computation, it calls *RESEND-RESULT* (Algorithm 9) in order to send back this result to the senders through a DAG of subquorums S_j^3 's, for $n + 1 \leq j \leq m + n$.

Algorithm 7 $REQUEST(R_1, R_2, i)$ $\triangleright r$ requests n senders through a DAG of subquorums, S_j^1 's, for $n + 1 \leq j \leq m + n$, to redo the computation.

- 1: r sets *REQ* to be a message consisting of R_1, R_2, i and r .
 - 2: r broadcasts *REQ* to all parties of quorum Q_{m+n} .
 - 3: All parties in Q_{m+n} calculate party, $q_{m+n} \in U_{m+n}$, of index $R_1[m, |U_{m+n}|]$ to be added to S_{m+n}^1 .
 - 4: The parties in Q_{m+n} send *REQ* to the parties of S_{m+n}^1 .
 - 5: **for** $j \leftarrow m + n, \dots, n + 2$ **do** \triangleright forwards *REQ* from output quorum to input quorums.
 - 6: **for all** $k : n + 1 \leq k < j$ and $(Q_k, Q_j) \in \text{Circuit}$ **do**
 - 7: All i parties in S_j^1 calculate party, $q_k \in U_k$, of index $R_1[k - n, |U_k|]$ to be added to S_k^1 .
 - 8: The parties in S_j^1 forward *REQ* and the IDs of all parties in S_k^1 to party q_k .
 - 9: Party q_k sends *REQ* to all the parties in S_k^1 .
 - 10: **end for**
 - 11: **end for**
 - 12: **for** $k \leftarrow n, \dots, 1$ **do** \triangleright Input quorums forward *REQ* to all senders.
 - 13: All i parties in the right neighboring subquorums of Q_k broadcast *REQ* to all parties in Q_k .
 - 14: All parties in Q_k send *REQ* to sender s_k .
 - 15: **end for**
-

Moreover, in *ELECT* (Q), the parties of quorum Q use MPC to elect an unmarked party uniformly at random from Q . Note that at any moment at least half of the unmarked parties in Q are good, thus the elected party is good with probability at least $1/2$. Finally, this election protocol runs in $O(1)$ time, and requires $O(\log^2 n)$ messages and $O(\log n)$ computational operations.

During *CHECK*, if any party receives inconsistent messages or fails to receive and verify any expected message in any round, it initiates a call to *UPDATE*.

2.4 UPDATE

When a computation is corrupted and *CHECK* detects this corruption, *UPDATE* is called. The *UPDATE* algorithm is described formally as Algorithm 10. When *UPDATE* starts, all parties in each quorum in the circuit are notified.

The main purpose of *UPDATE* is to 1) determine the location in which the corruption occurred; and 2) mark the parties that are in conflict.

To determine the location in which the corruption occurred, *UPDATE* calls *INVESTIGATE* (Algorithm 11) to investigate the current situation by letting each party involved in *COMPUTE*-

Algorithm 8 RECOMPUTE(R, i, r) \triangleright The n senders provide inputs to a DAG of subquorums, S_j^2 's, for $n + 1 \leq j \leq m + n$, in order to recompute producing a result, b_{m+n} , at r .

```

1: for each sender  $s_j$  that receives  $REQ$ ;  $1 \leq j \leq n$  and  $n + 1 \leq k, k' \leq m + n$  do  $\triangleright$  provides the inputs.
2:    $s_j$  broadcasts its input  $a_j, i$  and  $R$  to all parties in  $Q_j$ .
3:   All parties in  $Q_j$  calculate at most two right neighboring parties,  $q_k \in U_k$  and  $q_{k'} \in U_{k'}$ , of indices
    $R[k - n, |U_k|]$  and  $R[k' - n, |U_{k'}|]$  to be added to the right neighboring subquorums,  $S_k^2$  and  $S_{k'}^2$ .
4:   All parties in  $Q_j$  send  $a_j, i$  and  $R$  to the parties in  $S_k^2$  and the parties in  $S_{k'}^2$ .
5: end for
6: for  $j \leftarrow n + 1, \dots, m + n - 1$  do  $\triangleright$  recomputes
7:   for all  $k : j < k \leq m + n$  and  $(Q_j, Q_k) \in \text{Circuit}$  do
8:     if the new added party,  $q_j \in S_j^2$ , receives all its inputs then
9:        $q_j$  performs an operation on its inputs producing an output  $b_j$ .
10:       $q_j$  sends  $b_j, i$  and  $R$  to the parties of  $S_j^2$ .
11:      All  $i$  parties in  $S_j^2$  calculate party,  $q_k \in U_k$ , of index  $R[k - n, |U_k|]$  to be added to  $S_k^2$ .
12:      The parties in  $S_j^2$  send  $b_j, i, R$  and the IDs of all parties in  $S_k^2$  to party  $q_k$ .
13:    end if
14:  end for
15: end for
16: All  $i$  parties in  $S_{m+n}$  broadcast  $b_{m+n}, i$  and  $R$  to all parties in  $Q_{m+n}$ .
17: All parties in  $Q_{m+n}$  send  $b_{m+n}, i$  and  $R$  to party  $r$ .  $\triangleright r$  receives the result.

```

Algorithm 9 RESEND-RESULT(R, b_{m+n}, i) \triangleright Party r sends back the result of the computation, b_{m+n} , through a DAG of subquorums, S_j^3 's, to n senders, for $n + 1 \leq j \leq m + n$.

```

1:  $r$  sets  $RESULT$  to be a message consisting of  $R, b_{m+n}, i$  and  $r$ .
2:  $r$  broadcasts  $RESULT$  to all parties of quorum  $Q_{m+n}$ .
3: All parties in  $Q_{m+n}$  calculate party,  $q_{m+n} \in U_{m+n}$ , of index  $R[m, |U_{m+n}|]$  to be added to  $S_{m+n}^3$ .
4: The parties in  $Q_{m+n}$  send  $RESULT$  to the parties of  $S_{m+n}^3$ .
5: for  $j \leftarrow m + n, \dots, n + 2$  do  $\triangleright$  forwards result from output quorum to input quorums.
6:   for all  $k : n + 1 \leq k < j$  and  $(Q_k, Q_j) \in \text{Circuit}$  do
7:     All  $i$  parties in  $S_j^3$  calculate party,  $q_k \in U_k$ , of index  $R[k - n, |U_k|]$  to be added to  $S_k^3$ .
8:     The parties in  $S_j^3$  forward  $RESULT$  and the IDs of all parties in  $S_k^3$  to party  $q_k$ .
9:     Party  $q_k$  sends  $RESULT$  to all the parties in  $S_k^3$ .
10:  end for
11: end for
12: for  $k \leftarrow n, \dots, 1$  do  $\triangleright$  Input quorums forward the result to all senders.
13:   All  $i$  parties in the right neighboring subquorums of  $Q_k$  broadcast  $RESULT$  to all parties in  $Q_k$ .
14:   All parties in  $Q_k$  send  $RESULT$  to party  $s_k$ .
15: end for

```

Algorithm 10 UPDATE \triangleright Party $q' \in Q'$ calls $UPDATE$ after it claims that it detects a corruption.

```

1:  $q'$  broadcasts the fact that it calls  $UPDATE$  along with the messages it has received in this call to  $COMPUTE$ 
   to all parties in  $Q'$ .
2: The parties in  $Q'$  verify that  $q'$  received inconsistent messages before proceeding.
3:  $Q'$  notifies all quorums in the circuit via all-to-all communication that  $UPDATE$  is called.
4: INVESTIGATE  $\triangleright$  investigates all participants to determine corruption locations.
5: MARK-IN-CONFLICTS  $\triangleright$  marks the parties that are in conflict.

```

Algorithm 11 INVESTIGATE ▷ investigates parties to determine corruption locations.

```

1: for each party,  $q$ , involved in the last call to COMPUTE-CIRCUIT or CHECK do
2:    $q$  compiles all messages they have received (and from whom) and they have sent (and to whom) in the
   last call to COMPUTE-CIRCUIT or CHECK.
3:    $q$  broadcasts these messages to all parties in its quorum and neighboring quorums.
4: end for

```

CIRCUIT or *CHECK* broadcast all messages they have received or sent. Then, *UPDATE* calls *MARK-IN-CONFLICTS* (Algorithm 12) in order to mark the parties that are *in conflict*, where a pair of parties is in conflict if at least one of these parties broadcasted messages that conflict with the messages broadcasted by the other party in this pair. Note that each pair of parties that is in conflict has at least one bad party. Recall that if $(1/2 - \gamma)$ -fraction of parties in any quorum are marked, for any constant $\gamma > 0$, e.g., $\gamma = 0.01$, they are set unmarked. Also, for each pair of leaders that get marked, their quorums elect another pair of unmarked leaders.

Algorithm 12 MARK-IN-CONFLICTS ▷ marks the parties that are in conflict.

```

1: for each pair of parties,  $(q_x, q_y)$ , that is in conflict*, in quorums  $(Q_x, Q_y)$  do
2:   party  $q_y$  broadcasts a conflict message,  $\{q_x, q_y\}$ , to all parties in  $Q_y$ .
3:   each party in  $Q_y$  forwards  $\{q_x, q_y\}$  to all parties in  $Q_x$ .
4:   all parties in  $Q_x$  (or  $Q_y$ ) send  $\{q_x, q_y\}$  to the other quorums that has  $q_x$  (or  $q_y$ ).
5:   each quorum has  $q_x$  or  $q_y$  sends  $\{q_x, q_y\}$  to its neighboring quorums.
6: end for
7: for each party  $q$  that receives conflict message  $\{q_x, q_y\}$  do
8:    $q$  marks  $q_x$  and  $q_y$  in its marking table.
9: end for
10: if  $(1/2 - \gamma)$ -fraction of parties in any quorum have been marked, for  $\gamma = 0.01$  then
11:   each of these parties is set unmarked in all its quorums.
12:   each of these parties is set unmarked in all its neighboring quorums.
13: end if
14: for each pair of leaders,  $(q_x, q_y)$ , that is in conflict, in quorums  $(Q_x, Q_y)$  do
15:   ELECT( $Q_x$ ) and ELECT( $Q_y$ ) to elect a pair of unmarked leaders,  $(q'_x, q'_y)$ .
16:    $Q_x$  and  $Q_y$  notify their neighboring quorums with  $(q'_x, q'_y)$ .
17: end for

```

* A pair of parties, (q_x, q_y) , is *in conflict* if: 1) q_x was scheduled to send an output to q_y at some point in the last call to *COMPUTE-CIRCUIT* or *CHECK*; and 2) q_y does not receive an expected message from q_x in *INVESTIGATE*, or q_y receives a message in *INVESTIGATE* that is different than the message that it has received from q_x in the last call to *COMPUTE-CIRCUIT* or *CHECK*.

3 Analysis

In this section, we prove Theorem 1. All logarithms are base 2.

Definition 1. Rooted Directed Acyclic Graph (R-DAG) is a DAG in which, for a vertex u called the root and any other node v , there is at least one directed path from v to u .

Lemma 1. Given any R-DAG, of size n , in which each node has indegree of at most d and survives independently with probability at most p such that $0 < p \leq \frac{1}{d} - \epsilon$, for any constant $\epsilon > 0$, then

the probability of having a subgraph, rooted at some node, having only surviving nodes, of size $\Omega(\frac{\log n}{(1-pd)^2})$ is at most $1/2$.

Proof. This proof makes use of the following three propositions, but first we define some notations.

Given an R-DAG, $D(V, E)$, with size n and maximum indegree d , after each node survives independently with probability at most p such that $0 < p \leq \frac{1}{d} - \epsilon$, for any constant $\epsilon > 0$, we explore D to find a subgraph with only surviving nodes of size more than k , rooted at an arbitrary node v (assuming that node v survives).

Let $D'(v)$ be the maximum subgraph of surviving nodes, rooted at node v . Let each node in D have a status, which is either *inactive*, *active* or *neutral*. During the exploration process, the status of nodes is changed. A node x is inactive if $x \in D'(v)$ and its children are explored determining which one is in $D'(v)$. A node x is active if $x \in D'(v)$ but its children are not explored yet. A node x is neutral if it is neither active nor inactive, i.e., node x and its children are not explored yet.

The exploration process runs in at most $k > 0$ steps. Initially, we set an arbitrary surviving node, v , active and all other nodes neutral. At each step i , we choose an active node, w_i , in an arbitrary way, and we explore all its children. For all $(w_i, w'_i) \in E$ and w'_i survives and is neutral, we set w'_i active, otherwise w'_i remains as it is. Then, we set w_i inactive. Note that at any step, if there is no active node, the exploration process terminates. Now let d_i be the maximum number of children of node w_i for $1 \leq i \leq k$, i.e.,

$$d_i = \begin{cases} \deg(w_i) - 1 & \text{if } w_i \in V - \text{root}(D), \\ \deg(w_i) & \text{otherwise.} \end{cases}$$

where $\deg(w_i)$ is the degree of node w_i and $\text{root}(D)$ is the root node of D . For $1 \leq i \leq k$, let X_i be a non-negative random variable for the number of surviving neutral children of w_i , and let Y_i be a non-negative random variable for the number of surviving non-neutral children of w_i . Note that $Y_1 = 0$. So X_i follows a binomial distribution with parameters $(d_i - Y_i)$ and p , i.e., $X_i \sim \text{Bin}(d_i - Y_i, p)$. Let A_i be a non-negative random variable for the total number of active nodes after i steps, for $1 \leq i \leq k$.

Proposition 1. $A_i = \begin{cases} \sum_{j=1}^i X_j - (i-1) & \text{if } A_{i-1} > 0, \\ 0 & \text{otherwise.} \end{cases}$

Proof. Since the process starts initially with one active node v , $A_0 = 1$. Now we have two cases of A_{i-1} to compute A_i , $1 \leq i \leq k$:

Case 1 (process terminates before i steps): If $A_{i-1} = 0$, then $A_j = 0$ for $i \leq j \leq k$.

Case 2 (otherwise): If $A_{i-1} > 0$, then $A_i = A_{i-1} + X_i - 1$, where after exploring w_i , the total number of active nodes is the number of new active nodes (X_i) due to the exploration of w_i in addition to the total number of active nodes of previous steps (A_{i-1}) excluding w_i that becomes inactive at the end of step i . \square

Now let $|D'(v)|$ be the number of nodes in $D'(v)$.

Proposition 2. $\Pr(|D'(v)| > k) \leq \Pr(\sum_{i=1}^k X_i \geq k)$.

Proof. To prove this proposition, we first prove that $\Pr(|D'(v)| > k) \leq \Pr(A_k > 0)$. In order to do that, we prove that $|D'(v)| > k \implies A_k > 0$. If $|D'(v)| > k$, then the exploration process does not terminate before k steps. This implies that after k steps, there are k inactive nodes and at least one active node remains. This follows that $A_k > 0$. Thus, we have $\Pr(|D'(v)| > k) \leq \Pr(A_k > 0)$.

Now we prove that $\Pr(A_k > 0) \leq \Pr(\sum_{i=1}^k X_i - (k-1) > 0)$. To do that, we prove that $A_k > 0 \implies \sum_{i=1}^k X_i - (k-1) > 0$. If $A_k > 0$, then $A_j > 0$ for all $1 \leq j \leq k$.

By Proposition 1, we obtain that $\sum_{i=1}^j X_i - (j-1) > 0$ for all $1 \leq j \leq k$. This follows that $Pr(A_k > 0) \leq Pr(\sum_{i=1}^k X_i - (k-1) > 0)$.

Now, we obtain $Pr(|D'(v)| > k) \leq Pr(\sum_{i=1}^k X_i - (k-1) > 0)$, or equivalently, $Pr(|D'(v)| > k) \leq Pr(\sum_{i=1}^k X_i > k-1)$. Since k is a positive integer, we have $Pr(|D'(v)| > k) \leq Pr(\sum_{i=1}^k X_i \geq k)$. \square

Proposition 3. $Pr(\sum_{i=1}^k X_i \geq k) \leq e^{-\frac{(1-pd)^2 k}{1+pd}}$.

Proof. To prove this proposition, we first make use of stochastic dominance. For $1 \leq i \leq k$, let $X_i^+ \sim Bin(d, p)$, and let X_1^+, \dots, X_k^+ be independent random variables. We know that $Y_i \geq 0$ and $d_i \leq d$ for $1 \leq i \leq k$. By Theorem (1.1) part (a) of [15], for all $1 \leq i \leq k$, X_i^+ first-order stochastically dominates X_i , i.e., X_i^+ is stochastically larger than X_i . Hence, $\sum_{i=1}^k X_i^+$ is stochastically larger than $\sum_{i=1}^k X_i$. Thus, we have $Pr(\sum_{i=1}^k X_i \geq k) \leq Pr(\sum_{i=1}^k X_i^+ \geq k)$.

Now let $S_k = \sum_{i=1}^k X_i^+$. By Chernoff bounds, for $\delta > 0$, we obtain

$$Pr(S_k \geq (1+\delta)E(S_k)) \leq \left(\frac{e^\delta}{(1+\delta)^{(1+\delta)}} \right)^{E(S_k)} \leq e^{-\frac{\delta^2}{2+\delta} E(S_k)}.$$

We know that $S_k \sim Bin(kd, p)$. Thus, $E(S_k) = pdk$. Therefore, we have

$$Pr(S_k \geq (1+\delta)pdk) \leq e^{-\frac{\delta^2}{2+\delta} pdk}.$$

For $\delta = \frac{1-pd}{pd}$, we obtain $Pr(S_k \geq k) \leq e^{-\frac{(1-pd)^2 k}{1+pd}}$. \square

Now by Propositions 2 and 3, we have $Pr(|D'(v)| > k) \leq e^{-\frac{(1-pd)^2 k}{1+pd}}$. We know that node v survives with probability at most p . Thus, we obtain

$$Pr(|D'(v)| > k) \leq pe^{-\frac{(1-pd)^2 k}{1+pd}}.$$

Union bound over n nodes, then the probability that there exists a subgraph of D , rooted at one node, having only surviving nodes of size more than k is at most

$$nPr(|D'(v)| > k) \leq npe^{-\frac{(1-pd)^2 k}{1+pd}}.$$

Note that $npe^{-\frac{(1-pd)^2 k}{1+pd}} \leq 1/2$, when $k \geq \frac{1+pd}{(1-pd)^2 \log e} \log(2pn)$. Thus, the probability of having such a subgraph of size more than $\frac{1+pd}{(1-pd)^2 \log e} \log(2pn)$, or equivalently, $\Omega\left(\frac{\log n}{(1-pd)^2}\right)$, is at most $1/2$. \square

Corollary 1. For any R-DAG, of size n , the probability of having a subgraph, rooted at one node, having only surviving nodes, of size at least $n/2$ is $o(1)$.

Recall that in each round of *CHECK*, the new added parties are chosen uniformly at random. We define a *deception DAG*, D_i , as a subgraph of the DAG containing the new added parties, that has the following properties: 1) it has only bad parties, of which the adversary makes use to corrupt the computation; 2) it receives all its inputs, and each input is provided correct by at least one good party; and 3) each output it produces is provided to at least one good party, and at least one output is dropped or corrupted.

When the adversary corrupts at least one output of a deception DAG in any round, it has to keep doing such corruption in all subsequent deception DAGs; otherwise, the good parties that expect to receive certain messages from the deception DAG will receive inconsistent messages and so they call *UPDATE*. Even if the adversary keeps corrupting these outputs from a round to another, the corruption will be detected if the deception DAG shrinks to size zero at any round.

In the following facts, we show that 1) any deception DAG never expands in any direction from a round to another; and 2) it shrinks logarithmically from a round to another with probability at least $1/2$. This will imply that any deception DAG shrinks to size zero in $O(\log^* m)$ rounds with probability at least $1/2$.

Fact 2 *The deception DAG in any round is a supergraph of all subsequent deception DAGs.*

Let p be the probability of selecting an unmarked bad party uniformly at random in any quorum. Recall that the fraction of bad parties in any quorum is at most $1/4$, and at any time the fraction of unmarked parties in any quorum is at least $1/2 + \gamma$, for any constant $\gamma > 0$. Thus, the fraction of unmarked bad parties in any quorum is $\frac{1/2}{1+2\gamma}$, i.e., $p \leq \frac{1/2}{1+2\gamma}$.

Fact 3 *With probability at least $1/2$, any deception DAG, of size m , that is rooted at one party, shrinks to size zero in $O(\log^* m)$ rounds, where $0 < p \leq \frac{1/2}{1+2\gamma}$, for any constant $\gamma > 0$.*

Proof. Given a deception DAG, of size m , that is rooted at one party. By Fact 2, the deception DAG never expands over rounds. Let X_i be an indicator random variable, which is 1 if the deception DAG shrinks logarithmically from a round i to round $i + 1$; and 0 otherwise. Note that in each round, the receiver that is elected by the output quorum is good with probability at least $1/2$. By Lemma 1, the probability that the deception DAG shrinks logarithmically is at least $1/2$ given that the elected receiver is good. Since the probability of electing a good receiver is independent of the probability that the deception DAG shrinks logarithmically, $X_i = 1$ with probability at least $1/4$.

Now we show the required number of the X_i random variables to be 1 in order to shrink the deception DAG of size m to 0. Now fix p . Due to the constant factor, $c \leq \frac{3}{(1-pd)^2}$, of the logarithmic shrinking in Lemma 2, then after having $(\log^* m - 2)$ of X_i 's equal 1, the deception DAG of size m will shrink to size $2c(\log 2c)$. Moreover, Lemma 2 will not be applicable after the deception DAG shrinks to a constant size, $k \leq c \log k$. For such case, we simply make use of Corollary 1.

Thus, after having $(\log^* m - 2)$ of X_i 's equal 1, by Corollary 1, we further require at most $\log(2c(\log 2c))$, of X_i 's equal 1 to eventually shrink the deception DAG to size 0. This implies that we require at most $2 \log^* m$ of the X_i random variables to be 1.

Let $X = \sum_{i=1}^{16 \log^* m} X_i$. Then $E(X) = 4 \log^* m$, and since the X_i 's are independent, by Chernoff bounds,

$$\Pr(X < (1 - \delta)4 \log^* m) \leq \left(\frac{e^{-\delta}}{(1 - \delta)^{1+\delta}} \right)^{4 \log^* m}.$$

$$\text{For } \delta = \frac{1}{2} \text{ and } m > 2, \Pr(X < 2 \log^* m) \leq \left(\frac{e^{-\frac{1}{2}}}{(\frac{3}{2})^{\frac{3}{2}}} \right)^{4 \log^* m} < \frac{1}{2}. \quad \square$$

Lemma 2. *With probability at least $1/2$, any deception DAG, of size m , shrinks to size zero in $O(\log^* m)$ rounds, where $0 < p \leq \frac{1/2}{1+2\gamma}$, for any constant $\gamma > 0$.*

Proof. Let D be a deception DAG with multiple outputs and the adversary corrupts more than one output. Now let D_m be the maximum subgraph of D , that is rooted at one party of a corrupted output. Note that by definition any maximal subgraph that is rooted at one party of a corrupted output is a deception DAG. By Fact 2, each of these deception DAGs never expands over rounds. By Fact 3, each of these deception DAGs shrinks to size zero in a number of rounds at most the number of rounds that D_m shrinks to size zero.

Now we consider the case where the adversary maintains more than one deception DAG in the same round. By Fact 2 and the definition of deception DAG, the deception DAGs of the same round do not overlap, and they shrink independently. Also, by Fact 3, any maximum deception DAG, of one root, shrinks to size zero in a number of rounds that is at least the number of rounds that any other deception DAG, of one root, shrinks to size zero.

Therefore, for the adversary to maximize the expected number of rounds in *CHECK* is to consider the maximum deception DAG, of one root, in the first round of *CHECK*. We know that this maximum deception DAG has size at most m . By Fact 3, it shrinks to size zero in $O(\log^* m)$. \square

The next lemma shows that *CHECK* catches corruptions with probability $\geq 1/2$.

Lemma 3. *Assume some party selected uniformly at random in the last call to COMPUTE-CIRCUIT has corrupted a computation. Then when the algorithm CHECK is called, with probability at least $1/2$, some party will call UPDATE.*

Lemma 4. *If some party selected uniformly at random in the last call to COMPUTE-CIRCUIT or CHECK has corrupted a computation, then UPDATE will identify a pair of neighboring quorums Q and Q' such that at least one pair of parties in these quorums is in conflict and at least one party in such pair is bad.*

The next lemma bounds the number of calls to *UPDATE* before all bad parties are marked.

Lemma 5. *UPDATE is called $O(t)$ times before all bad parties are marked.*

4 Conclusion and Future Work

We have presented algorithms for reliable multiparty computations. These algorithms can significantly reduce message cost and number of computational operations to be asymptotically optimal. The price we pay for this improvement is the possibility of computation corruption. In particular, if there are $t \leq (\frac{1}{4} - \epsilon)n$ bad parties, for any constant $\epsilon > 0$, our algorithm allows $O(t(\log^* m)^2)$ computations to be corrupted in expectation.

Many problems remain. First, it seems unlikely that the smallest number of corruptions allowable by an attack-resistant algorithm with optimal message complexity is $O(t(\log^* m)^2)$. Can we improve this to $O(t)$ or else prove a non-trivial lower bound? Second, we allow the inputs of parties to reveal. Can we maintain the privacy of these inputs? Third, we assume a partially synchronous communication model, which is crucial for our *CHECK* algorithm to detect computation corruptions over rounds. Can we extend this algorithm to fit for asynchronous computations?

References

1. Baruch Awerbuch and Christian Scheideler. Towards a scalable and robust dht. *Theory of Computing Systems*, 45(2):234–260, 2009.
2. Donald Beaver. Efficient multiparty protocols using circuit randomization. In Joan Feigenbaum, editor, *Advances in Cryptology CRYPTO 91*, volume 576 of *Lecture Notes in Computer Science*, pages 420–432. Springer Berlin Heidelberg, 1992.
3. Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, pages 1–10, New York, NY, USA, 1988. ACM.
4. Iching Boman, Jared Saia, Chaouki Abdallah, and Edl Schamiloglu. Brief announcement: Self-healing algorithms for reconfigurable networks. In *Stabilization, Safety, and Security of Distributed Systems*, volume 4280 of *Lecture Notes in Computer Science*, pages 563–565. Springer Berlin Heidelberg, 2006.
5. A. Das Sarma and A. Trehan. Edge-preserving self-healing: keeping network backbones densely connected. In *IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 226–231, 2012.

6. Amos Fiat and Jared Saia. Censorship resistant peer-to-peer networks. *Theory of Computing*, 3(1):1–23, 2007.
7. Amos Fiat, Jared Saia, and Maxwell Young. Making chord robust to byzantine attacks. In *Proceedings of the 13th annual European Symposium on Algorithms, ESA'05*, pages 803–814. Springer-Verlag, 2005.
8. T. Frisanco. Optimal spare capacity design for various protection switching methods in atm networks. In *Proceedings of International Conference on Communications*, volume 1 of *ICC'97*, pages 293–298, 1997.
9. Thomas P. Hayes, Jared Saia, and Amitabh Trehan. The forgiving graph: a distributed data structure for low stretch under adversarial attack. In *Proceedings of the 28th ACM symposium on Principles of Distributed Computing*, PODC '09, pages 121–130, 2009.
10. Tom Hayes, Navin Rustagi, Jared Saia, and Amitabh Trehan. The forgiving tree: a self-healing distributed data structure. In *Proceedings of the 27th ACM symposium on Principles of Distributed Computing*, PODC '08, pages 203–212, 2008.
11. Kirsten Hildrum and John Kubiawicz. Asymptotically efficient approaches to fault-tolerance in peer-to-peer networks. In *Distributed Computing*, volume 2848 of *Lecture Notes in Computer Science*, pages 321–336. Springer Berlin Heidelberg, 2003.
12. R.R. Iraschko, M.H. MacGregor, and W.D. Grover. Optimal capacity placement for path restoration in stm or atm mesh-survivable networks. *IEEE/ACM Transactions on Networking*, 6(3):325–336, 1998.
13. Aniket Kate and Ian Goldberg. Distributed key generation for the internet. In *Proceedings of the 29th IEEE International Conference on Distributed Computing Systems, ICDCS '09*, pages 119–128, 2009.
14. Valerie King, Steven Lonargan, Jared Saia, and Amitabh Trehan. Load balanced scalable byzantine agreement through quorum building, with full information. In *Proceedings of the 12th International Conference on Distributed Computing and Networking*, pages 203–214. Springer-Verlag, 2011.
15. Achim Klenke and Lutz Mattner. Stochastic ordering of classical discrete distributions. *Advances in Applied Probability*, 42(2):392 – 410, 2010.
16. Jeffrey Knockel, George Saad, and Jared Saia. Self-healing of byzantine faults. In *SSS*, pages 98–112, 2013.
17. K. Murakami and H.S. Kim. Comparative study on restoration schemes of survivable atm networks. In *16th Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 1 of *INFOCOM '97*, pages 345–352, 1997.
18. Moni Naor and Udi Wieder. A simple fault tolerant distributed hash table. In *2nd International Workshop on Peer-to-Peer Systems*, pages 88–97, 2003.
19. Gopal Pandurangan and Amitabh Trehan. Xheal: localized self-healing using expanders. In *Proceedings of the 30th annual ACM SIGACT-SIGOPS symposium on Principles of Distributed Computing*, PODC '11, pages 301–310, 2011.
20. Manoj Prabhakaran and Amit Sahai. *Secure Multi-Party Computation*, volume 10. IOS Press, 2013.
21. T. Rabin and M. Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, STOC '89, pages 73–85, 1989.
22. Jared Saia and Amitabh Trehan. Picking up the pieces: Self-healing in reconfigurable networks. In *IEEE International Symposium on Parallel and Distributed Processing*, IPDPS'08, pages 1–12, 2008.
23. Christian Scheideler. How to spread adversarial nodes? rotate! In *Proceedings of the 37th annual ACM Symposium on Theory of Computing*, STOC '05, pages 704–713, 2005.
24. B. Van Caenegem, N. Wauters, and P. Demeester. Spare capacity assignment for different restoration strategies in mesh survivable networks. In *Proceedings of International Conference on Communications*, volume 1 of *ICC'97*, pages 288–292, 1997.
25. Yijun Xiong and L.G. Mason. Restoration strategies and spare capacity requirements in self-healing atm networks. *IEEE/ACM Transactions on Networking*, 7(1):98–110, 1999.
26. Andrew C. Yao. Protocols for secure computations. In *Proceedings of the 23rd Annual Symposium on Foundations of Computer Science*, SFCS '82, pages 160–164, 1982.

A Appendix - Deferred Proofs

Fact 2. The deception DAG in any round is a supergraph of all subsequent deception DAGs.

Proof. We know by definition that the deception DAG is bordered by the good parties that provide the inputs to the DAG, and the good parties that receive the outputs from the DAG.

In each round, all parties of each subquorum in round i send their outputs to the new added party in the next subquorum. Thus, the good parties that provide the correct inputs to the deception DAG of round i , will provide the correct inputs to the deception DAGs in all subsequent rounds. Moreover, in each round, the new added party in each subquorum forwards its output to all parties in its subquorum. Note that each good party has previously received a message, it verifies this message with all subsequent messages it receives, and if it receives inconsistent messages or fails to receive an expected message, then it calls *UPDATE*.

Therefore, all good parties that border a deception DAG in any round will border all subsequent deception DAGs. \square

Lemma 3. Assume some party selected uniformly at random in the last call to *COMPUTE-CIRCUIT* has corrupted a computation. Then when *CHECK* is called, with probability at least $1/2$, some party will call *UPDATE*.

Proof. Recall that in *REQUEST*, in each round for $O(\log^* m)$ rounds, a new party r is elected by the output quorum to send a request to the senders in order to recompute. By Lemma 2, this request is sent reliably with probability at least $1/2$ to all input quorums. Note that each request has a round number. Thus, at any round, if any good party in any input quorum receives a request of round number i and has not received $(i - 1)$ requests of proper round numbers, then it calls *UPDATE*.

If all input quorums receive all requests properly in all $O(\log^* m)$ rounds, then *RECOMPUTE* must be called properly $O(\log^* m)$ times by all input quorums. By Lemma 2, the result is computed and sent reliably to the output quorum with probability at least $1/2$.

Recall that in *RECOMPUTE*, the round number i is forwarded with the computation results from the senders to the output quorum. Thus, at any round, if any good party in the output quorum receives a result with a round number i and has not received $(i - 1)$ results with proper rounds numbers, then it calls *UPDATE*.

Now, if all parties in the output quorum receive all results properly in all $O(\log^* m)$ rounds, then *RESEND-RESULT* must be called $O(\log^* m)$ times by the output quorum. By Lemma 2, the result of the computation is sent back reliably to all senders with probability at least $1/2$. Thus, the probability that *CHECK* succeeds in finding a corruption and calling *UPDATE* is at least $1/2$. \square

Lemma 4. If some party selected uniformly at random in the last call to *COMPUTE-CIRCUIT* or *CHECK* has corrupted a computation, then *UPDATE* will identify a pair of neighboring quorums Q and Q' such that at least one pair of parties in these quorums is in conflict and at least one party in such pair is bad.

Proof. First, we show that if a pair of parties x and y is in conflict, then at least one of them is bad. Assume not. Then both x and y are good. This implies that party x would have truthfully reported what it received and sent; any result that x has computed would have been sent directly to y ; and y would have truthfully reported what it received from x . But this is a contradiction, since for x and y to be in conflict, y must have reported that it received from x something different than what x reported sending.

Now consider the case where a selected unmarked bad leader corrupted the computation in the last call to *COMPUTE-CIRCUIT*. By Lemma 3, with probability at least $1/2$, some party, $q' \in Q'$, will call *UPDATE*. Recall that in *UPDATE* q' broadcasts all messages it has received to all parties in Q' . These parties verify if q' received inconsistent messages before proceeding.

In *UPDATE*, we know that each party, $q \in Q$, participated in the last call to *COMPUTE* broadcasts what it has received and sent to all parties in Q . Thus, all parties of Q verify the correctness of q 's computation. Thus, if the corruption occurs due to an incorrect computation made by a bad party, this corruption will be detected and all parties will know that this party is bad.

Now if all parties compute correctly and *CHECK* detects a corruption, then we show that there is some pair of parties will be in conflict. Assume this is not the case. Thus, by the definition of corruption, there must be a deception DAG, in which all inputs are provided correct and an output is corrupted at party q' . Then each pair of parties, $(q_j, q_k) \in (Q_j, Q_k)$, in the deception DAG that is rooted at q' , is not in conflict, for $n + 1 \leq j < k \leq m + n$. Thus, we have that 1) this DAG received all its inputs correct; 2) all parties compute correctly; and 3) no pair of parties is in conflict. This implies that it must be the case that q' received the correct output. But if this is the case, then q' that initially called *UPDATE* would have received no inconsistent messages. This is a contradiction since in such a case, this party would have been unsuccessful in trying to initiate a call to *UPDATE*. Thus, *UPDATE* will find two parties that are in conflict, and at least one of them will be bad. \square

Lemma 5. *UPDATE* is called $O(t)$ times before all bad parties are marked.

Proof. By Lemma 4, if a corruption occurred in the last call to *COMPUTE-CIRCUIT*, and it is caught by *CHECK*, then *UPDATE* is called. *UPDATE* identifies at least one pair of parties that is in conflict, and each of such pairs has at least one bad party.

Now let g be the number of marked good parties, and let b be the number of marked bad parties. Also let $f(b, g) = b - (\frac{p}{1-p})g$. Since $0 < p \leq \frac{1/2}{1+2\gamma}$, for any constant $\gamma > 0$, $0 < \frac{p}{1-p} \leq \frac{1}{1+4\gamma}$.

For each corruption caught, at least one bad party is marked, and so $f(b, g)$ increases by at least $\frac{1-2p}{1-p}$ since b increases by at least 1 and g increases by at most 1. When $(1/2 - \gamma)$ -fraction of parties in any quorum Q get unmarked, for any constant $\gamma > 0$, $f(b, g)$ further increases by at least 0 since b decreases by at most $p|Q|(1/2 - \gamma)$ and g decreases by at least $(1-p)|Q|(1/2 - \gamma)$. Hence, $f(b, g)$ is monotonically increasing by at least $\frac{1-2p}{1-p}$ for each corruption caught. When all bad parties are marked, $f(b, g) \leq t$. Therefore, after at most $(\frac{1-p}{1-2p})t$, or at most $(1 + \frac{1}{2\gamma})t$, calls to *UPDATE*, all bad parties are marked. \square

Proof of Theorem 1. We first show the message cost, the number of operations and the latency of our algorithms. By Lemma 5, the number of calling *UPDATE* is at most $O(t)$. Thus, the resource cost of all calls to *UPDATE* is bounded as the number of calls to *COMPUTE-CIRCUIT* grows large. Therefore, for the amortized cost, we consider only the cost of the calls to *COMPUTE-CIRCUIT* and *CHECK*.

When a computation is performed through a circuit of m gates with a circuit depth ℓ , *COMPUTE-CIRCUIT* has message cost $O(m + n \log n)$, number of operations $O(m + n \log n)$ and latency $O(\ell)$. *CHECK* has message cost $O((m + n \log n)(\log^* m)^2)$, number of operations $O((m + n \log n) \log^* m)$ and latency $O(\ell \log^* m)$, but *CHECK* is called only with probability $1/(\log^* m)^2$. Hence, the call to *CHECK* has an amortized expected message cost $O(m + n \log n)$, amortized computational operations $O(\frac{m+n \log n}{\log^* m})$ and an amortized expected latency $O(\ell / \log^* m)$.

In particular, if we call *COMPUTE* \mathcal{L} times, then the expected total number of messages sent will be $O(\mathcal{L}(m + n \log n) + t(m \log^2 n))$ with expected total number of computational operations $O(\mathcal{L}(m + n \log n) + t(m \log n \log^* m))$ and latency $O(\ell(\mathcal{L} + t))$. This is true since *UPDATE* is called $O(t)$ times and each call to *UPDATE* has message cost $O(m \log^2 n)$ with computational operations $O(m \log n \log^* m)$ and latency $O(\ell)$.

Recall that by Lemma 5, the number of times *CHECK* must catch corruptions before all bad parties are marked is $O(t)$. In addition, if a bad party caused a corruption during a call to *COMPUTE-CIRCUIT*, then by Lemmas 3 and 4, with probability at least $1/2$, *CHECK* will catch it. As a consequence, it will call *UPDATE*, which marks the parties that are in conflict. *UPDATE* is thus called with probability $1/(\log^* m)^2$, so the expected total number of corruptions is $O(t(\log^* m)^2)$.